Redes Neurais de Grafos versus Modelos Gravitacionais Estruturais: Uma Comparação Metodológica para Impactos no Comércio Agrícola Decorrentes de Rejeições de Micotoxinas na União Europeia.

Graph Neural Networks versus Structural Gravity Models: A Methodological Comparison for Agricultural Trade Impacts from EU Mycotoxin Rejections.

Leonardo Ferraz
ESALQ/USP. Piracicaba, SP, Brasil
Heloisa Lee Burnquist
ESALQ/USP. Piracicaba, SP, Brasil
Jéssica Jesus Teixeira
UFV. Viçosa, MG, Brasil
Michelle Márcia Viana Martins
UFV. Viçosa, MG, Brasil
Fernando Elias Corrêa
USP. Piracicaba, SP, Brasil

GT01. Mercados agrícolas e comércio exterior

Resumo: Este estudo compara Redes Neurais de Grafos (GNNs) com modelos gravitacionais estruturais na análise de impacto no comércio agrícola causado por rejeições de micotoxinas na União Europeia. A aplicação de padrões SPS rigorosos pela UE resulta em rejeições de fronteira que redirecionam exportações para mercados alternativos, com implicações significativas para a segurança alimentar global. Utilizando dados bilaterais de comércio do UN COMTRADE (2008–2023) ao nível HS-6 para commodities suscetíveis a micotoxinas e notificações do sistema RASFF, implementamos modelos GNN estrutural e flexível baseados em passagem de mensagens, incorporando a intensidade de rejeições por micotoxinas do RASFF como atributo de aresta. A metodologia de regressão simbólica é empregada para interpretar as representações aprendidas pela GNN, investigando se o modelo descobre de forma autônoma estruturas funcionais análogas aos termos de resistência multilateral da teoria gravitacional. Os resultados indicam que a GNN aprende representações economicamente coerentes do comércio agrícola, com embeddings de países que refletem estruturas regionais e assimetrias de acessibilidade bilateral consistentes com a literatura gravitacional. A regressão simbólica redescobre estruturas funcionais compatíveis com os termos de resistência multilateral, incluindo interações entre tamanho econômico e agregados ponderados pela acessibilidade, mesmo sem imposição de restrições teóricas. Esses achados estabelecem a base metodológica para a aplicação futura de simulações contrafactuais de desvio de comércio induzido por rejeições de micotoxinas.

Palavras-chave: redes neurais de grafos, modelo gravitacional, termos de resistência multilaterais, micotoxinas, regressão simbólica.

Abstract: This study compares Graph Neural Networks (GNNs) with structural gravity models in analyzing agricultural trade impacts from EU mycotoxin rejections. The EU's enforcement of stringent SPS standards results in border rejections that redirect exports toward alternative markets, with significant implications for global food safety. Using bilateral trade data from UN COMTRADE (2008–2023) at the HS-6 level for mycotoxin-susceptible commodities and RASFF notifications, we implement both Structural and Flexible GNN models, incorporating RASFF mycotoxin rejection intensity as an edge feature. Symbolic regression is employed to interpret the GNN's learned representations, investigating whether the model autonomously discovers functional structures analogous to multilateral resistance terms from gravity theory. Results indicate that the GNN learns economically coherent representations, with country embeddings reflecting regional structures and bilateral accessibility asymmetries consistent with the gravity literature. Symbolic regression rediscovers functional structures compatible with multilateral resistance terms, including interactions between economic size and accessibility-weighted aggregates, even without imposed theoretical restrictions. These findings lay the methodological groundwork for future counterfactual simulations of trade diversion induced by mycotoxin rejections.

Keywords: graph neural networks, gravity model, multilateral resistance terms, mycotoxins, symbolic regression.

1. Introduction

The gravity model of international trade, in its modern structural form derived from Anderson and van Wincoop (2003), has become the dominant framework for empirical trade analysis. Its theoretical foundation — based on constant elasticity of substitution preferences, Armington differentiation, and market-clearing conditions — yields a reduced-form expression in which bilateral trade flows are determined by the economic sizes of the trading partners, their bilateral trade costs, and multilateral resistance terms (MRTs) that capture each country's average trade costs with all of its partners. Fixed effects effectively control for MRTs, but at the cost of absorbing relevant variation without revealing underlying mechanisms, precluding out-of-sample predictions, and generating computational challenges in disaggregated data, what might be considered as limitations that motivate the search for alternative approaches.

A particularly promising direction was introduced by Verstyuk and Douglas (2022), who applied Graph Neural Networks (GNNs) to gravity models of international trade. Using a large panel of aggregate bilateral trade data (68 years, 181 countries), they demonstrated that a flexible GNN-based model outperforms both the structural gravity model and the fixed-effects specification. Using symbolic regression to interpret trained GNNs, they rediscovered gravity-consistent multilateral resistance structures alongside novel functional forms, with their flexible model (R^2 of 0.68) outperforming both structural (R^2 of 0.53) and fixed-effects benchmarks (R^2 of 0.62), while supply-side market characteristics consistently explained nearly twice the variance of their demand-side counterparts.

The present study applies this framework to agricultural trade in commodities susceptible to mycotoxin contamination. The EU's Rapid Alert System for Food and Feed (RASFF) documents thousands of border rejections each year, and this network structure — where a rejection at the EU border potentially reshuffles trade patterns across dozens of alternative destinations — exemplifies the multilateral interdependence that GNNs are designed to model. The data environment (HS-6 product level, sixteen-year panel, narrow commodity scope) is substantially more challenging than the aggregate flows used in the original study, offering a rigorous test of the methodology's applicability.

We implement both Structural and Flexible GNN models following Verstyuk and Douglas (2022), incorporating RASFF mycotoxin rejection intensity as an edge feature. The structural gravity specification serves here as a theoretical reference framework, not as an estimated competing model: it provides the economic vocabulary — exporter capabilities (S_j), importer market characteristics (M_i), bilateral accessibility (φ_{ij}), multilateral resistance —

against which we assess whether the GNN's learned representations are economically coherent rather than atheoretical artifacts.

The central contribution lies in applying symbolic regression (PySR) to the trained Flexible GNN. The research question is: when a GNN learns to model agricultural trade flows in mycotoxin-susceptible commodities without explicit theoretical constraints, does symbolic regression reveal that its learned representations admit an economically interpretable structure grounded in the theoretical foundations of the gravity framework — and, if so, to what extent do the discovered functional forms relate to multilateral resistance concepts from the Anderson–van Wincoop (2003) tradition?

This question is simultaneously methodological and conceptual. On the methodological side, it tests whether the Verstyuk and Douglas (2022) findings — obtained on aggregate trade across all products over nearly seven decades — extend to a narrower commodity range, and more policy-relevant data environment. On the conceptual side, the discovered functional forms, results from the symbolic regression, may provide insights into how trade in agricultural commodity markets differ from aggregate trade in the way multilateral resistance operates, with potential implications for the design of trade policy analysis tools. The explicit modeling of trade diversion — using the trained models to simulate counterfactual reallocation of trade flows following regulatory shocks such as mycotoxin rejections — is discussed as a direction for future research.

2. Related Literature

2.1. The EU's Mycotoxin Regulatory Framework and Border Rejections

The EU enforces some of the world's strictest mycotoxin maximum residue levels (MRLs), grounded in its precautionary approach and WTO SPS Agreement prerogatives that allow standards exceeding Codex Alimentarius benchmarks when scientifically justified. For aflatoxins in groundnuts destined for direct human consumption, for instance, the EU limit of 4 μg/kg is less than half the Codex standard of 10 μg/kg and considerably more restrictive than thresholds applied across Asia, Africa, and Latin America (WU, 2004).

RASFF data document each border rejection event — product, hazard type and concentration, origin country, and notifying member state — revealing systematic patterns whereby certain origins and commodities generate disproportionate mycotoxin notifications, with frequencies fluctuating alongside climatic conditions, regulatory threshold changes, and enforcement intensity. These patterns constitute the empirical foundation of the study, providing a source of trade cost variation encoded as an edge feature in the bilateral accessibility network — capturing a regulatory dimension of market access that standard gravity covariates do not reflect.

2.2. Agricultural Trade Costs as a Network Phenomenon

In mycotoxin-susceptible commodity markets, regulatory barriers introduce multilateral trade cost heterogeneity that standard gravity covariates do not capture: when EU rejection rates rise for a given exporter, cost adjustments propagate through the entire network, conditioned by the simultaneous repositioning of competing suppliers. Therefore, trade cost adjustments depend not only on the affected exporter's relationships with alternative buyers but on how competing exporters — Argentina, the United States, India, and China, in the case of peanuts — are simultaneously repositioned across the same markets.

This network interdependence — theoretically acknowledged by structural gravity through MRTs but left unresolved by fixed-effects specifications — motivates the GNN approach, whose message-passing architecture is explicitly designed to learn representations that incorporate broader neighborhood structure rather than treating bilateral relationships in isolation. If rejections operate primarily through flow redirection rather than bilateral accessibility compression, the message-passing mechanism is better equipped to capture these network-wide diversion effects than any bilateral parameter alone.

2.3. Gravity Models in Agricultural Trade Analysis

The structural gravity specification follows the canonical form derived from Anderson and van Wincoop (2003) and is implemented empirically in the manner recommended by Head and Mayer (2014). In empirical practice, MRTs are typically controlled for by including exporter-time and importer-time fixed effects, as recommended by Baldwin and Taglioni (2007). The PPML estimator proposed by Santos Silva and Tenreyro (2006) handles zero trade flows and is robust to heteroskedasticity. Fixed effects effectively control for MRTs, but at the cost of absorbing relevant variation without revealing underlying mechanisms, precluding out-of-sample predictions, and generating computational challenges in disaggregated data.

We note that the PPML gravity specification is not estimated in the present study; rather, its theoretical decomposition into exporter capabilities, importer market characteristics, and bilateral accessibility provides the interpretive framework against which we evaluate whether the GNN's learned representations are economically grounded.

2.4. Machine Learning in Trade Economics

GNNs remain underexplored in trade analysis, though early work shows their capacity to capture network effects and, under certain conditions, produce node embeddings interpretable as learned multilateral resistance terms. Gradient boosting and random forest methods have been used for trade flow prediction and anomaly detection (Starosta de Waldemar, 2020). Monken et al. (2021) demonstrated the potential of GNN architectures for trade flow prediction, and Jiang (2023) provided a theoretical treatment showing that under certain conditions, GNN node embeddings can be interpreted as learned representations of multilateral resistance terms.

Building on Verstyuk and Douglas (2022), this study applies GNNs to SPS-driven trade diversion — a context well-suited to network analysis given its multilateral dynamics — using the structural gravity model not as a competing estimator but as an interpretive benchmark to assess whether the GNN's learned representations are economically coherent rather than atheoretical artifacts.

3. Methodology

3.1. Data Sources and Construction

Trade data come from COMTRADE via CEPII's BACI reconciled dataset, covering bilateral flows at the HS-6 level for mycotoxin-susceptible commodities — maize and maize derivatives (HS chapter 1005), soybeans and soybean products (HS chapter 1201), groundnuts and peanut products (HS chapters 1202 and 2008.11), and cocoa and cocoa products (HS chapter 18) — over a sixteen-year panel (2008–2023), with zero flows explicitly retained.

SPS shocks are identified from the European Commission's RASFF database, using mycotoxin notifications for 2008–2023. From these, we construct a rejection intensity variable — rejections per million USD of bilateral trade — normalized by trade scale for each origin country, product category, and year, which enters the model as an edge feature alongside standard gravity covariates.

Standard gravity covariates are drawn from the CEPII GeoDist and Gravity databases (Head and Mayer, 2014). The bilateral trade cost proxy variables used as edge features in the φ network comprise sixteen variables: a regional trade agreement indicator (RTA), contiguity, common official language, common ethnic language, common colonizer post-1945, colonial relationship post-1945, log population-weighted distance in kilometers, hegemon indicators for both origin and destination, any historical colonial relationship, colonial dependency status, a common religion index, FTA/WTO membership, GATT/WTO membership indicators for both origin and destination, and the RASFF rejection intensity variable. Time-varying country characteristics (GDP, GDP per capita, population) are sourced from the World Bank's World Development Indicators. Following Verstyuk and Douglas (2022), expenditure X_i is defined as GDP plus total imports minus total exports, while production Y_j is represented by GDP.

The sample covers 65 countries, of which 60 participate as importers and 50 as exporters, spanning all major regions of the global agricultural trade network. South America contributes nine countries including Brazil, Argentina, Chile, Colombia, and Peru; the Middle East and North Africa contributes twelve countries including Turkey, Egypt, Saudi Arabia, the United Arab Emirates, Israel, and Morocco; Sub-Saharan Africa contributes ten countries including Nigeria, South Africa, Kenya, Ghana, Côte d'Ivoire, and Ethiopia; Southeast Asia contributes seven countries including Indonesia, Thailand, Vietnam, and Malaysia; South Asia contributes six countries including India, Pakistan, and Bangladesh; Europe contributes six non-EU countries including Switzerland, Russia, and Ukraine; Central America and the Caribbean contributes five countries; East Asia contributes four countries including China, Japan, and South Korea; North America contributes the United States, Canada, and Mexico; Oceania contributes Australia and New Zealand; and Central Asia is represented by Kazakhstan. Fifteen countries appear exclusively as importers (including Angola, Bangladesh, Hong Kong, Jordan, Saudi Arabia, and several Gulf states), while five appear exclusively as exporters (Bolivia, Cameroon, Ethiopia, Myanmar, and Paraguay), reflecting the directional asymmetry of agricultural commodity trade in these product categories.

3.2. Structural Gravity Model

Although we do not estimate a PPML fixed effects regression in this study, a description of the structural gravity specification is essential for two reasons: it establishes the theoretical benchmark against which we interpret the GNN's learned representations, and it clarifies the structural restrictions — on supply capabilities, demand characteristics, and multilateral resistance terms — that the GNN is free to discover or depart from.

The general form of the gravity equation for trade is:

$$X_{ijp_t} = \exp(\beta_{i_t} + \beta_{j_t} + \beta_{p_t} + \gamma \ln d_{ij} + \delta \, \text{Rejection}_{jp_t} + \delta' \, \text{Rejection}_{jp,_{t-1}} + \zeta \mathbf{Z}_{ij}) \times \xi_{ijp_t}$$

where X_{ijp_t} denotes the trade value from exporter j to importer i in product p, at time t; β_{i_t}, β_{j_t}, and β_{p_t} are importer-time, exporter-time, and product-time fixed effects that jointly control for multilateral resistance terms (MRTs), time-varying country characteristics, and global product-specific demand and supply shocks; d_{ij} is bilateral distance; $Rejection_{jp_t}$ is the contemporaneous mycotoxin rejection intensity for exporter j and product p; Z_{ij} is a vector of bilateral gravity covariates (contiguity, common language, colonial ties, RTA membership); and ξ_{ijp_t} is the error term.

The structural gravity specification serves here as a theoretical reference framework, not as an estimated competing model. The gravity equation's decomposition into exporter capabilities (S_j), importer market characteristics (M_i), and bilateral accessibility (φ_{ij}) provides the conceptual vocabulary through which we interpret the GNN's outputs, while the multilateral resistance terms provide the functional forms against which we evaluate the symbolic regression results.

3.3. Graph Neural Network Architecture

The trade system is represented as a directed annual graph G = (V, E), where nodes are countries and directed edges are observed bilateral flows, each carrying feature vectors encoding country characteristics and bilateral trade determinants — including the RASFF rejection intensity variable. Through message-passing layers, each node iteratively aggregates neighborhood information and updates its representation according to the update rule at layer l:

$$h_u^{(l+1)} = \mathrm{UPDATE}\left(h_u^{(l)}, \mathrm{AGGREGATE}(\{ h_v^{(l)} : v \in N(u) \}) \right)$$

where $h_u^{(l)}$ is the embedding of node u at layer l, N(u) is the set of neighbors of u, UPDATE is a learned function that produces the new node embedding, and AGGREGATE is a permutation-invariant function that combines the neighbor embeddings. The predicted trade flow $\ln \hat{X}$ from i to j at time t is decomposed as:

$$\ln \hat{X}_{ij,t} = \ln G + \ln M_{i,t} + \ln S_{j,t} + \ln \varphi_{ij,t} + \xi_{ij,t}$$

where G is a learnable scaling constant, $M_{i,t}$ captures importing market characteristics, $S_{j,t}$ captures exporter capabilities, and $\varphi_{ij,t}$ represents bilateral and neighborhood-level trade cost factors. The key innovation is that $M_{i,t}$, $S_{j,t}$, and $\varphi_{ij,t}$ are not pre-specified as fixed effects or parametric functions of observables but are learned end-to-end from the data.

Two architectural properties — exchangeability and locality — equip GNNs to capture multilateral resistance-like effects. Following Verstyuk and Douglas (2022), a single message-passing layer suffices given the density of agricultural trade networks, while also avoiding over-smoothing. Their interaction produces adaptive multilateral resistance: learned country representations that reflect network position and competitive exposure, and that generalize to unobserved country-time combinations — a capability fixed effects fundamentally lack.

Our implementation in PyTorch Geometric comprises seven neural network components and one learnable scalar, totaling 4,858 parameters. The bilateral accessibility network f^{φ} maps the sixteen-dimensional edge feature vector to a scalar $\varphi_{ij} \in [0, 1]$ through a feedforward architecture with two hidden layers (50 and 20 neurons), ReLU activations, dropout regularization ($p = 0.3$), and a sigmoid output function. Four inner functions ($f^M_{\varphi Y}$, $f^M_{\varphi X}$, $f^S_{\varphi X}$, and $f^S_{\varphi Y}$) each take as input the pair (φ_{ij}, GDP variable) and produce a scalar message through a single hidden layer of 20 neurons with layer normalization, ReLU activation, and dropout ($p = 0.2$). Two outer functions (f^M and f^S) receive the country's own economic characteristic together with the two aggregated message sums, producing $\ln(M_i)$ and $\ln(S_j)$ through two hidden layers (50 and 20 neurons) with ReLU activations and dropout ($p = 0.3$). The Structural Model baseline uses the same φ network but replaces the learned outer functions with theoretically constrained multilateral resistance terms Φ_i and Ω_j, computed iteratively following equations (7)–(8) of Verstyuk and Douglas (2022).

3.4. Symbolic Regression for Interpretability

To address ML interpretability concerns, we apply symbolic regression (SR) following Verstyuk and Douglas (2022) and Cranmer et al. (2020). After training the GNN, we generate synthetic input-output pairs from the trained model and feed them to the Python Symbolic Regression (PySR) algorithm, which searches millions of candidate algebraic expressions — combining addition, multiplication, division, exponentiation, and logarithms — selecting those on the Pareto frontier of accuracy and parsimony, via a genetic algorithm.

The central question this procedure answers is whether the GNN, trained without explicit multilateral resistance terms, implicitly discovers functional forms resembling the theoretical MRT expressions from structural gravity. If so, this would validate the gravity model's theoretical structure even at the disaggregated, regulatory-barrier-shaped commodity level; if not, it would suggest that standard MRT specifications inadequately characterize multilateral resistance in agricultural markets, motivating theoretical refinement.

3.5. Training and Evaluation

The GNN minimizes MSE on log trade flows via Adam optimizer, using a temporal split — 2008–2019 for training, 2020–2021 for validation and early stopping, and 2022–2023 as the held-out test set. Tuned hyperparameters include embedding dimension (16, 32, 64, or 128), hidden units, learning rate, and weight decay, with implementation via PyTorch Geometric. Performance is evaluated across predictive accuracy (RMSE and MAE, reported in- and out-of-sample), directional accuracy (AUC-ROC), computational efficiency, and interpretability assessed through the economic plausibility and parsimony of symbolic regression outputs.

4. Results and Discussion

4.1. Model Estimation and Predictive Performance

Both models were estimated by minimizing the MSE criterion via stochastic gradient descent, following Verstyuk and Douglas (2022). The Structural model was trained for 1,000 epochs, while the Flexible GNN employed early stopping after approximately 130 epochs based on validation loss. Figure 1 presents the training curves for both models.

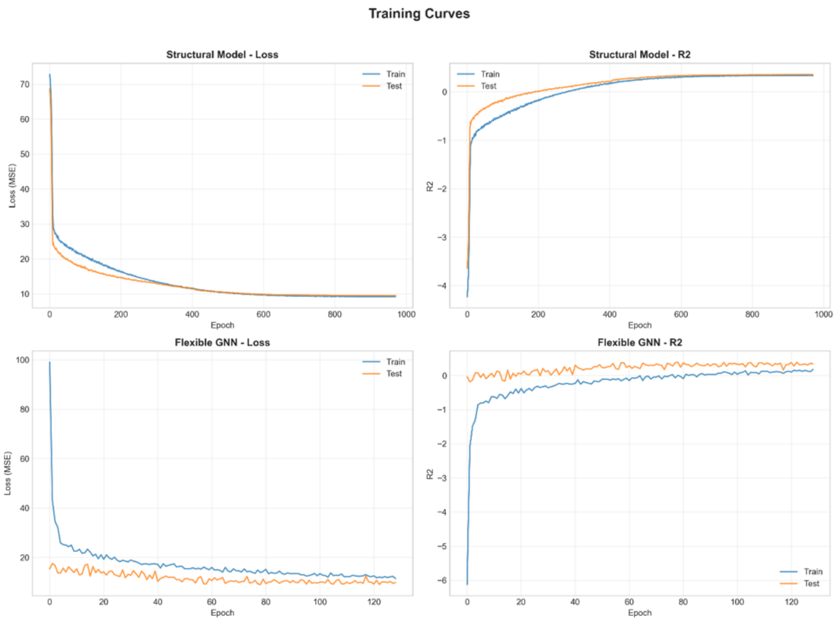

Figure 1. Training curves for the Structural and Flexible GNN models.

Top row: Structural model loss (left) and R^2 (right) over 1,000 epochs. Bottom row: Flexible GNN loss (left) and R^2 (right) over 130 epochs with early stopping. **Source:** Authors' elaboration.

The structural model converges slowly (~600 epochs) with tight train-test loss alignment, indicating no overfitting but limited capacity. The flexible GNN converges rapidly (~80 epochs), with early stopping preventing overfitting. Table 1 reports comparative performance against the Verstyuk and Douglas (2022) benchmarks.

Table 1. Model performance comparison: out-of-sample R^2

Model	R^2 IS	R^2 OOS	Parameters	Epochs
Structural (Ours)	0.36	0.36	1,892	1,000
Flexible GNN (Ours)	0.39	0.39	4,858	~130
Structural (V&D 2022)	0.53	0.51	382	1,000
Flexible (V&D 2022)	0.68	0.68	63,704	1,000

Source: Authors' elaboration. V&D 2022 values from Verstyuk and Douglas (2022), Table 1

Our models yield R^2 values of 0.36 (Structural) and 0.39 (Flexible GNN), lower than Verstyuk and Douglas (2022)'s 0.53 and 0.68. This gap can be explained by HS-6 disaggregation, a narrow and volatile commodity scope, and a shorter panel (2008–2023 vs. 68 years). Nonetheless, the relative ordering is preserved, the Flexible GNN outperforms the structural baseline, and the close alignment between in-sample and out-of-sample R^2 confirms the absence of overfitting in both specifications.

4.2. Supply and Demand Market Characteristics

Figure 2 displays the GNN's learned representations of importer market characteristics and exporter supply potential — generated through message-passing for each country-year — revealing the underlying structure of the international agricultural trade network.

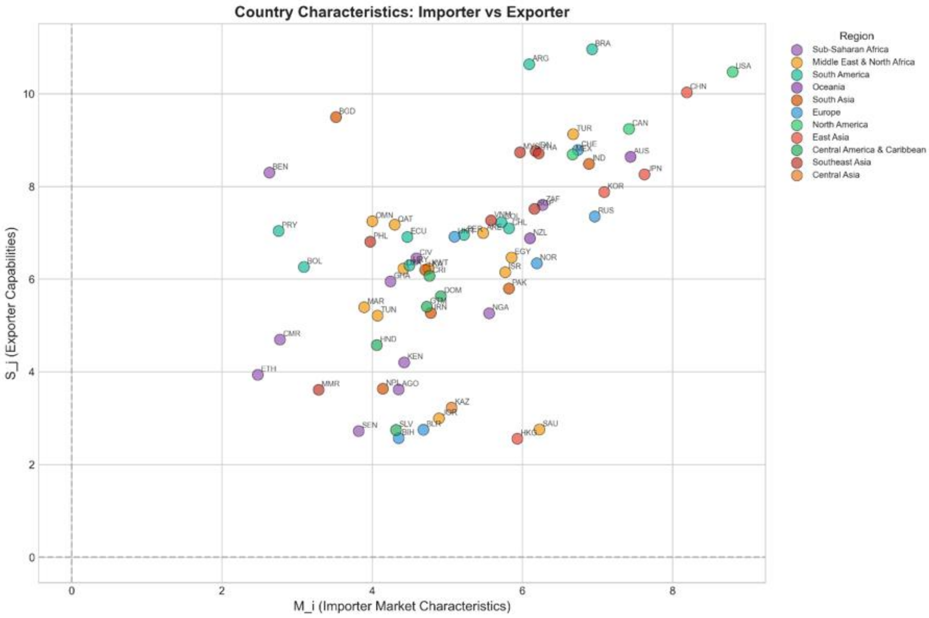

Figure 2. Country characteristics: importer market size (M_i) versus exporter capabilities (S_j) as learned by the Flexible GNN. Countries colored by geographic region. **Source:** Authors' elaboration

Several patterns are noteworthy. The major agricultural exporters — Brazil, the United States, Argentina, China, and Canada — occupy the upper-right quadrant, exhibiting both high exporter capabilities and substantial import market size. Brazil stands out with the highest S_j value in the sample, consistent with its position as the world's leading exporter of soybeans, the second-largest exporter of maize, and a major exporter of cocoa and peanut products. The United States achieves the highest M_i value, reflecting its large domestic consumption market. Countries in Sub-Saharan Africa and South Asia cluster in the lower-left quadrant, consistent with their marginal role in global mycotoxin-susceptible commodity trade. The positive correlation between M_i and S_j across countries is consistent with the home market effect documented in the trade literature.

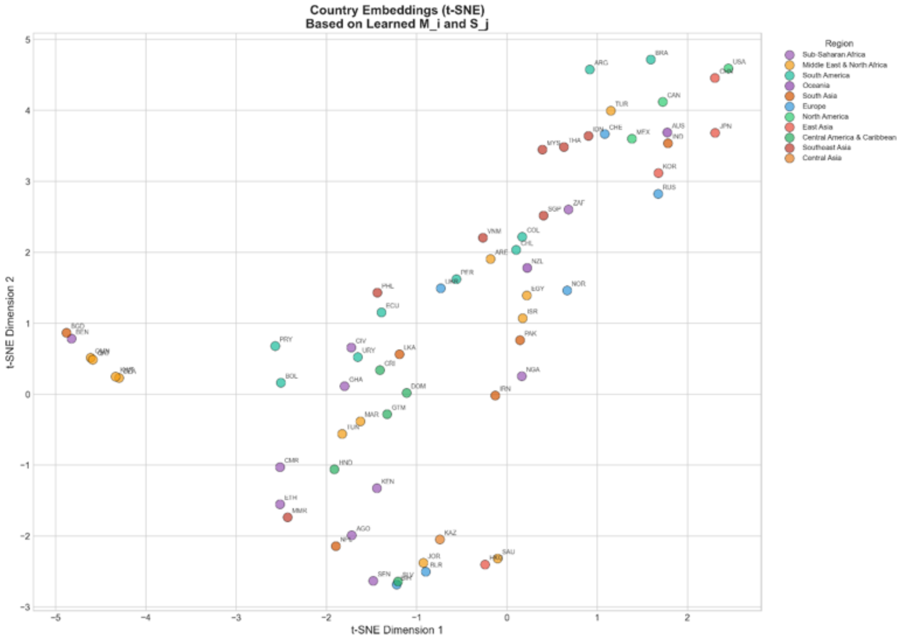

Figure 3. t-SNE visualization of country embeddings based on learned M_i and S_j characteristics. Countries colored by geographic region. **Source:** Authors' elaboration.

What the t-SNE projection reveals (Figure 3) is that this latent space possesses an economically interpretable geometry: European, South American, East Asian, and Sub-Saharan African countries occupy distinct regions, forming coherent regional clusters, despite no geographic information beyond bilateral distance having been provided as input. This indicates that the model, in optimizing its prediction of trade flows, internalized patterns of economic and institutional similarity that, in practice, organize themselves along regional lines — geographic proximity, trade agreements, specialization patterns — emerging spontaneously from the network's structure rather than from any explicit supervision. This latent space is the mathematical space where the learned country representations reside; its analysis via t-SNE is a way of auditing what the model has effectively "learned" about the structure of global agricultural trade.

4.3. Bilateral Accessibility

The bilateral accessibility parameter φ_{ij}, estimated by the auxiliary neural network from 16 trade cost proxy variables, captures the ease of market access between each country pair. Figure 4 presents the accessibility matrix for the 15 largest trading countries in the sample.

Figure 4. Bilateral accessibility matrix (φ_{ij}) for the top 15 countries by trade volume. Values range from 0 (no accessibility) to 1 (full accessibility). Rows represent importers; columns represent exporters. **Source:** Authors' elaboration.

The heatmap reveals several important patterns. First, bilateral accessibility is markedly asymmetric: φ_{ij} ≠ φ_{ji} for most country pairs. This asymmetry is consistent with Verstyuk and Douglas (2022) and reflects directional differences in regulatory barriers, tariff structures, and institutional trade costs. Second, China as an exporter achieves very high accessibility (φ ≈ 1.0) to most destinations, while as an importer from several partners the values are lower, reflecting the well-documented asymmetry of Chinese trade policy. Third, geographically proximate and economically integrated country pairs consistently exhibit high accessibility values, validating that the φ network captures the expected geographic and institutional determinants of trade costs.

4.4. Effect of Mycotoxin Rejections on Trade

Figure 6 presents the central empirical question: the effect of RASFF mycotoxin rejections on bilateral accessibility and trade flows.

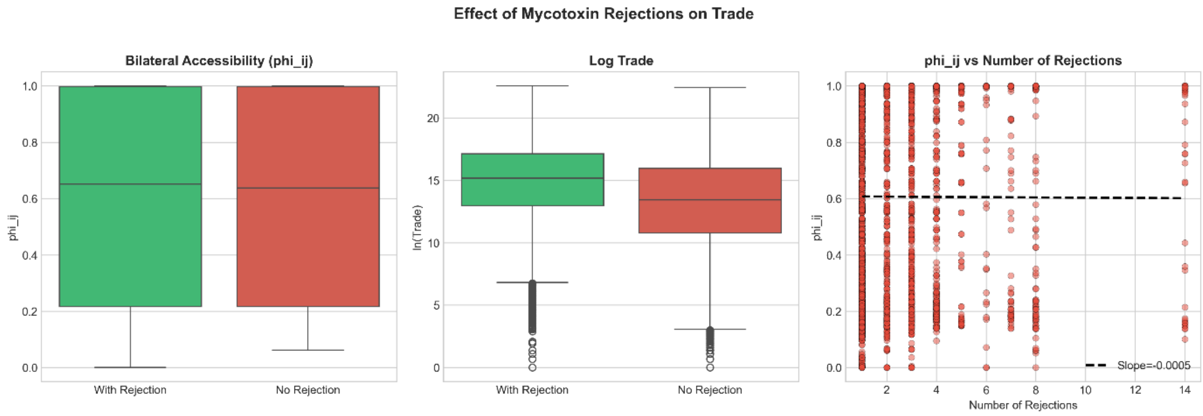

Figure 6. Effect of mycotoxin rejections on trade. Left: distribution of φ_{ij} by rejection status. Center: distribution of ln(Trade) by rejection status. Right: bilateral accessibility versus number of rejections, with linear trend (slope = −0.0005). **Source:** Authors' elaboration.

RASFF rejections introduce greater dispersion in bilateral market accessibility without shifting median values, indicating heterogeneous rather than uniformly negative effects across country pairs. In addition, the marginally higher trade volumes observed among rejected pairs reflect a selection effect — rejections occur only where trade is already active and substantial.

The near-zero slope (−0.0005) between rejection intensity and bilateral accessibility suggests that RASFF rejections exert negligible direct effects on structural trade costs, which are dominated by factors such as distance, language, and institutional quality. Rejections likely operate instead through trade diversion — reshaping flows toward third-country destinations rather than directly compressing bilateral accessibility.

This interpretation reinforces the case for the GNN approach: if rejections operate primarily through flow redirection rather than bilateral accessibility compression, the message-passing mechanism — which aggregates information across the full neighborhood structure — is better equipped to capture these network-wide diversion effects than any bilateral parameter alone.

4.5. Interpreting the GNN via Symbolic Regression

Following Verstyuk and Douglas (2022) and Cranmer et al. (2020), we applied PySR symbolic regression separately to each component of the trained GNN, searching for parsimonious closed-form expressions via a genetic algorithm optimized along the accuracy–complexity Pareto frontier. Results are reported in Table 2.

Table 2. Symbolic regression results for Flexible GNN components

Component	Arguments (z)	Symbolic form $\tilde{f}(z, \vartheta)$	Notes
$f^M_{\varphi Y}$	$\hat{\varphi}_{ik}$, Y_k	phi × ((Y_k × 0.234) − 0.780)	Inner, φ×Y
$f^S_{\varphi Y}$	$\hat{\varphi}_{ik}$, Y_k	Y_k × 0.390 × (phi − 0.386)	Inner, φ×Y
$f^M_{\varphi X}$	$\hat{\varphi}_{kj}$, X_k	(phi + X_k) × 0.170	Inner, φ×X
$f^S_{\varphi X}$	$\hat{\varphi}_{kj}$, X_k	((X_k + phi) × −0.147) + 0.255	Inner, φ×X
f^M	X_i, $\Sigma f_{\varphi Y}$, $\Sigma f_{\varphi X}$	((X_i + 3.066)/0.661) + |−1.686 − (X_i + 0.729) × $\Sigma f_{\varphi Y}$ × 0.064|	Outer, demand
f^S	Y_j, $\Sigma f_{\varphi X}$, $\Sigma f_{\varphi Y}$	($\Sigma f_{\varphi X}$ × 1.505 + exp(Y_j) − $\Sigma f_{\varphi Y}$) × 0.300 + 1.902	Outer, supply
V&D inner φY	φ, Y	$\vartheta_0 + (\vartheta_{10} + \vartheta_{11}z_1)/(\vartheta_{20} + \vartheta_{21}z_2)$	Benchmark
V&D outer M	X, ΣφY, ΣφX	$\vartheta_1 \ln((\vartheta_{10}+\vartheta_{11}z_1)/(\vartheta_{20}+\vartheta_{21}z_2)) + \vartheta_3 z_3$	Benchmark
V&D outer S	Y, ΣφX, ΣφY	$\vartheta_1 \ln((\vartheta_{10}+\vartheta_{11}z_1)/(\vartheta_{30}+\vartheta_{31}z_3)) + \vartheta_2 z_2$	Benchmark

Notes: V&D rows reproduce the unrestricted SR results from Verstyuk and Douglas (2022), Table 4.

The inner functions discovered by PySR exhibit predominantly multiplicative relationships between bilateral accessibility (φ) and the economic size variables (Y_k, X_k). For the φY components, both $f^M_{\varphi Y}$ and $f^S_{\varphi Y}$ take the form of Y_k multiplied by a function of φ: $f^M_{\varphi Y} = \varphi \times (0.234\, Y_k - 0.780)$ and $f^S_{\varphi Y} = 0.390\, Y_k \times (\varphi - 0.386)$. These multiplicative structures are economically intuitive, implying that a partner's production contributes to aggregate market access only to the extent modulated by bilateral accessibility — a relationship at the heart of the multilateral resistance concept. The $f^S_{\varphi Y}$ expression additionally introduces a threshold effect: only partners with accessibility above approximately 0.39 contribute positively. The φX components adopt simpler additive forms: $f^M_{\varphi X} = 0.170 \times (\varphi + X_k)$ and $f^S_{\varphi X} = -0.147 \times (X_k + \varphi) + 0.255$.

The outer functions reveal more complex and economically richer structures. For supply-side capabilities f^S, PySR discovered $(1.505 \times \Sigma f_{\varphi X} + \exp(Y_j) - \Sigma f_{\varphi Y}) \times 0.300 + 1.902$. Three features are noteworthy. First, the exporter's production Y_j enters through an exponential transformation, paralleling Verstyuk and Douglas's finding that logarithmic transformations emerged endogenously. Second, accessibility-weighted aggregate expenditure ($\Sigma f_{\varphi X}$) enters with a positive coefficient while accessibility-weighted aggregate production ($\Sigma f_{\varphi Y}$) enters with a negative sign — capturing a competitive pressure dynamic whereby an exporter's capabilities increase with access to large consumer markets but decrease when its importers have abundant alternative supply sources. Third, the overall elasticity with respect to Y_j is close to unity, consistent with near-unitary elasticities in the gravity literature.

For demand-side characteristics f^M, the discovered expression is $(X_i + 3.066)/0.661 + |-1.686 - (X_i + 0.729) \times \Sigma f_{\varphi Y} \times 0.064|$. The dominant term is a linear rescaling of importer expenditure X_i, which alone explains approximately 90 percent of the variance. The second term introduces a multiplicative interaction between X_i and the accessibility-weighted production of trading partners ($\Sigma f_{\varphi Y}$) — the closest analogue to a classical multilateral resistance term found in our analysis.

The principal convergence with Verstyuk and Douglas (2022) is that both analyses independently discover structures in which a country's trade characteristics depend not only on its own economic size but also on accessibility-weighted aggregates of partner characteristics — the defining feature of MRTs in the structural gravity framework. The principal divergence lies in functional forms: whereas Verstyuk and Douglas found logarithmic ratio structures, our expressions favor linear and exponential forms. This plausibly reflects that disaggregated agricultural trade over 16 years occupies a narrower region of the input space where linear approximations are locally accurate. Despite this difference, the economic content is preserved: the GNN independently learns that bilateral trade depends on the relative — not absolute — market position of each country, which is the foundational insight of Anderson and van Wincoop (2003).

4.6. Discussion

The results jointly address the study's central question across four dimensions. On predictive performance, the Flexible GNN modestly but consistently outperforms the Structural model ($R^2 = 0.39$ vs. 0.36), with robust in- and out-of-sample generalization. On multilateral resistance, the learned country embeddings replicate economically meaningful heterogeneity — regional clustering, interpretable M and S orderings — functioning as effective substitutes for explicit fixed effects. On bilateral accessibility, the φ_{ij} heatmap confirms economic geography determinants and pronounced asymmetry ($\varphi_{ij} \neq \varphi_{ji}$), with direct implications for trade policy. On mycotoxin rejections, the weak direct relationship between RASFF rejections and φ suggests that trade diversion operates primarily through network-level flow reallocation rather than bilateral accessibility compression — the channel the GNN's message-passing is best equipped to capture.

The symbolic regression results reinforce this interpretation: the competitive pressure term in the supply-side function captures how rejections alter the landscape for all network-connected exporters, while the demand-side interaction implies that countries with limited supply connectivity are most exposed to sudden inflows of diverted products. This asymmetry concentrates the food safety consequences of EU mycotoxin rejections in precisely those markets with the weakest regulatory infrastructure and fewest alternative suppliers — reinforcing existing vulnerabilities in the global food safety architecture.

5. Final Remarks

The results support three main findings. First, the Flexible GNN outperforms the Structural Model in both in-sample and out-of-sample fit ($R^2 = 0.39$ versus 0.36), confirming that the unconstrained message-passing architecture captures trade network effects beyond what the theoretically restricted specification allows, even in a narrow commodity context. Second, the GNN's learned country representations exhibit economically coherent patterns: major agricultural exporters such as Brazil, the United States, and Argentina are correctly positioned, regional clustering emerges without explicit geographic node features, and bilateral accessibility φ_{ij} displays the asymmetry documented in the original study. Third, and most importantly, symbolic regression reveals that the GNN independently discovers functional structures that admit interpretation in terms of multilateral resistance. In particular, the inner functions learn that a partner's production contributes to aggregate market access only when modulated by bilateral accessibility, and the supply-side expression introduces a threshold below which weakly connected partners are effectively filtered out — demonstrating that the gravity model's theoretical insight extends to disaggregated agricultural commodity markets, even when no theoretical structure is imposed.

Several limitations should be acknowledged. The model does not estimate causal effects of mycotoxin rejections on trade flows; identification as a causal trade cost shock would require an instrumental variable strategy beyond the scope of this study. The single message-passing layer limits the model's capacity to capture higher-order network effects. The absence of a PPML fixed-effects benchmark estimated on the same dataset means that the performance comparison is made against the Structural GNN rather than the current workhorse econometric specification. Notwithstanding these limitations, the analysis of mycotoxin rejections yields a finding with direct methodological implications: RASFF rejections introduce greater dispersion in bilateral accessibility without shifting its central tendency, and the near-zero slope between rejection intensity and structural trade costs suggests that their effects operate primarily through network-wide flow redirection rather than bilateral accessibility compression — a mechanism that the GNN's message-passing architecture is inherently better equipped to capture than any bilateral parameter alone.

The network interdependence revealed by these results — theoretically acknowledged by structural gravity through MRTs but left unresolved by fixed-effects specifications that control for it without modeling it — points toward two directions for future work that emerge naturally. The first and most immediate is the use of the trained models for counterfactual trade diversion simulations — embedding the functional forms discovered by symbolic regression in a general equilibrium adjustment loop to simulate how trade flows would reallocate following specific policy scenarios, such as the removal of mycotoxin rejections for a given origin country, harmonization of regulatory standards to the Codex Alimentarius level, or strengthening enforcement capacity in alternative destination markets. The second direction involves extending the analysis to cross-product linkages — modeling how mycotoxin contamination in maize affects soybean meal trade through animal feed supply chains — exploiting the GNN's capacity to model multiple interconnected commodity networks within a unified graph structure.

References

ANDERSON, J. E.; VAN WINCOOP, E. Gravity with gravitas: a solution to the border puzzle. **American Economic Review**, Washington, v. 93, n. 1, p. 170–192, 2003.

BALDWIN, R.; TAGLIONI, D. Trade effects of the euro: a comparison of estimators. **Journal of Economic Integration**, Sejong, v. 22, n. 4, p. 780–818, 2007.

CORREIA, S.; GUIMARÃES, P.; ZYLKIN, T. Fast Poisson estimation with high-dimensional fixed effects. **The Stata Journal**, College Station, v. 20, n. 1, p. 95–115, 2020.

CRANMER, M. Interpretable machine learning for science with PySR and SymbolicRegression.jl. **arXiv preprint**, arXiv:2305.01582, 2023. Disponível em: https://arxiv.org/abs/2305.01582. Acesso em: 15 mar. 2026.

CRANMER, M. *et al.* Discovering symbolic models from deep learning with inductive biases. *In*: ADVANCES IN NEURAL INFORMATION PROCESSING SYSTEMS, 33., 2020. **Proceedings** [...]. [S.l.]: NeurIPS, 2020.

EUROPEAN COMMISSION. **RASFF — Rapid Alert System for Food and Feed**. Brussels: European Commission, 2025. Disponível em: https://food.ec.europa.eu/safety/rasff-food-and-feed-safety-alerts_en. Acesso em: 15 mar. 2026.

HEAD, K.; MAYER, T. Gravity equations: workhorse, toolkit, and cookbook. *In*: GOPINATH, G.; HELPMAN, E.; ROGOFF, K. (ed.). **Handbook of International Economics**. Amsterdam: Elsevier, 2014. v. 4, p. 131–195.

JIANG, X. Graph neural networks for trade flow prediction: connections to the structural gravity model. **Working Paper**, 2023.

MONKEN, A. *et al.* Graph neural networks for modeling causality in international trade. *In*: NEURIPS WORKSHOP ON CAUSAL INFERENCE, 2021. **Proceedings** [...]. [S.l.]: NeurIPS, 2021.

SANTOS SILVA, J. M. C.; TENREYRO, S. The log of gravity. **The Review of Economics and Statistics**, Cambridge, v. 88, n. 4, p. 641–658, 2006.

STAROSTA DE WALDEMAR, F. Machine learning in trade analysis: gradient boosting and random forest applications. **Working Paper**, 2020.

UNITED NATIONS STATISTICS DIVISION. **UN COMTRADE Database**. New York: United Nations, 2025. Disponível em: https://comtrade.un.org/. Acesso em: 15 mar. 2026.

VERSTYUK, S.; DOUGLAS, M. R. Machine learning the gravity equation for international trade. **SSRN Working Paper**, n. 4053795, 2022. Disponível em: https://ssrn.com/abstract=4053795. Acesso em: 15 mar. 2026.

WU, F. Mycotoxin risk assessment for the purpose of setting international regulatory standards. **Environmental Science & Technology**, Washington, v. 38, n. 15, p. 4049–4055, 2004.